

MEMORANDUM

RM-4816-PR

MARCH 1967

417872

NEURONAL SPIKE TRAINS AND STOCHASTIC POINT PROCESSES

Donald H. Perkel, George L. Gerstein and George P. Moore

This research is supported by the United States Air Force under Project RAND—Contract No. F44620-67-C-0045—monitored by the Directorate of Operational Requirements and Development Plans, Deputy Chief of Staff, Research and Development, Hq USAF. Views or conclusions contained in this Memorandum should not be interpreted as representing the official opinion or policy of the United States Air Force.

DISTRIBUTION STATEMENT

Distribution of this document is unlimited.

The **RAND** Corporation

1700 MAIN ST. • SANTA MONICA • CALIFORNIA • 90406

PREFACE

The work reported here results from a collaboration between experimental neurophysiologists and mathematical and computer specialists, concerned with the statistical analysis and interpretation of trains of impulses obtained from individual nerve cells. Of the authors, Donald H. Perkel is at The RAND Corporation, and is a consulting member of the Brain Research Institute, University of California at Los Angeles; George L. Gerstein is with the Department of Biophysics (Johnson Research Foundation) and the Department of Physiology, School of Medicine, University of Pennsylvania; George P. Moore is with the Departments of Physiology and Engineering and the Brain Research Institute, University of California at Los Angeles, and is a consultant to The RAND Corporation.

This research has been supported by U.S. Air Force Project RAND, by the National Institutes of Health, through grants GM-09608, NB-06372, and NB-05606, and by the National Science Foundation, through grant G-21497.

Mathematical studies and computer modeling of neural systems are a part of the Project RAND research directed toward the better understanding of information processing in nervous systems.

The contents of this Memorandum have been accepted for publication in the Biophysical Journal, and are scheduled to appear in Vol. 7, No. 4 (July 1967).

SUMMARY

The train of impulses ("spikes") produced by a nerve cell is, in a growing class of neurophysiological experiments, subjected to statistical treatment in which the time intervals between spikes play an essential rôle. The statistical analysis of spike trains is developed here in terms of the underlying theory of stochastic point processes, i.e., of stochastic processes whose realizations may be described as series of point events occurring in time, separated by random intervals. It is intended to describe the principal techniques available in a fairly systematic fashion and to relate them to the underlying theory. Most of the computational techniques discussed have been reported in the literature; some are new.

For single stationary spike trains, several orders of complexity of statistical treatment are outlined; the most important distinction is that made between those statistical measures that depend in an essential way on the serial order of interspike intervals and those that are order-independent. The interrelations among the several types of calculations are shown, and an attempt is made to ameliorate the current nomenclatural confusion in this field. The application and interpretation of these statistical results are discussed, and some potential difficulties are outlined. Next, the related analysis is discussed for experiments in which a brief isolated stimulus is presented repeatedly.

The analysis of two simultaneously observed spike trains is discussed next. The first statistical question that arises is whether the observed trains are independent; statistical techniques for testing the independence hypothesis are developed around the notion that, under the null hypothesis, the times of spike occurrence in one train represent random instants in time with respect to the other train. If the null hypothesis is rejected--if dependence is attributed to the trains--the subsequent problems are those of characterizing the nature and source of the observed dependencies. These are discussed, largely by means of illustrations drawn from computer simulations of interacting neurons. The combination of repetitive stimulation and simultaneous recording from two (or more) neurons is shown to give more detailed clues as to possible interactions among the monitored neurons; the theory that is developed is illustrated by an application to experimental data from auditory neurons.


Finally, the effects of nonstationarity--e.g., long-term changes in firing rate--on the various statistical measures are discussed. The severity of the effects of rate changes on single spike-train statistics depends strongly on the "inherent" variability of the interspike intervals; regularly firing ("pacemaker") neurons are much more sensitive to accelerations and decelerations than are more irregularly firing cells. For two-train comparisons,

inherent difficulties that arise for pacemaker-like neurons even when stationary are exacerbated by shared rate changes. For nonpacemaker cells, however, the detection and measurement of dependency is significantly impaired only if the nonstationarities are fairly severe, and hence readily apparent in the individual spike-train data.

[The page contains extremely faint, illegible text, likely bleed-through from the reverse side of the document.]


ACKNOWLEDGEMENTS

The authors thank Dr. José P. Segundo and Professor Theodore H. Bullock for continued encouragement, discussions, and criticism. We are also very grateful to Dr. P. A. W. Lewis for valuable criticisms and suggestions.



CONTENTS

PREFACE.....	iii
SUMMARY.....	v
ACKNOWLEDGMENTS.....	ix
INTRODUCTION.....	1
Motivation.....	1
Basic Terminology and Scope.....	4
THE SINGLE SPIKE TRAIN.....	8
Stochastic Point Processes: Basic Nomenclature..	8
Order-independent Statistical Measures.....	10
Order-dependent Statistical Measures.....	24
Description, Prediction, and Information.....	54
SINGLE SPIKE TRAINS IN THE PRESENCE OF STIMULATION....	57
TWO SIMULTANEOUS SPIKE TRAINS.....	60
The Problem of Functional Relationship.....	60
Independence: Consequences of the Null Hypothesis	61
Application to Experimental Data.....	69
Types of Dependence.....	76
TWO SPIKE TRAINS IN THE PRESENCE OF STIMULATION.....	86
Functional Effects of Stimulation.....	86
Prediction of the Cross Correlation.....	91
Application to Experimental Data.....	96
THE PROBLEM OF NONSTATIONARITY.....	104
Basic Concepts.....	104
Measurement of Rate Variations.....	107
Effects of Trends.....	109
REFERENCES.....	134



NEURONAL SPIKE TRAINS AND STOCHASTIC POINT PROCESSES

INTRODUCTION

Motivation

A neuronal spike train is the sequence of nerve impulses, or action potentials, produced by a neuron, typically observed over a relatively long period of time. The analysis of spike trains has been of increasing interest to neurophysiologists in recent years, stimulated, no doubt, by wide availability of automatic data-processing equipment. Spike-train analysis differs from "classical" electrophysiological methods in that the raw data of interest are not precise voltage measurements, but rather precise measurements of times of occurrence of events. From these essentially temporal data, statistical descriptions are obtained of the output behavior of neurons, from which inferences may be made, in turn, with regard to certain specific types of basic neurophysiological questions. It is not our intention to pursue the questions of interpretation and inference in the present paper (for a recent review, see Moore, Perkel & Segundo, 1966), but rather to present in systematic fashion the statistical techniques of spike-train analysis, and in particular to point out some of the relevant mathematical assumptions and relationships underlying the computational techniques.

Methodological considerations, however, must not be isolated from inferential and interpretational questions; we have tried to discuss the computational techniques within the larger context of neurophysiological investigation. Our principal working assumptions, as stated in the study mentioned above (Moore, Perkel & Segundo, 1966), are "(a) that there is an enormous wealth of information about the structure and function of the nervous system which can be derived from careful study of the detailed timings of spike events; (b) that analysis of these signals can shed light on mechanisms of spike production within the observed cell, on the presynaptic input to the cell, and on the mechanisms by which the latter is transformed into a postsynaptic output; and (c) that observation of multiple units can reveal details of interconnections and functional interactions ... [and] the appreciation that neuronal processes at all levels involve a probabilistic element which must be adequately incorporated if quantitative hypotheses or models of neuronal functions are to be valid. Finally, it is held that only the more detailed analyses of spike timings are appropriate to any quantitative theory of information processing by the nervous system."

Spike-train analysis is applied at several different levels of interpretation, and it is the level of interpretation that largely dictates the choice and depth of statistical techniques for processing of the data. At

one level, the statistical measures of the spike train provide a relatively concise characterization of the output of the neuron, which may be used for description, comparison, and classification of nerve cells. At another level, spike-train statistics of a neuron may afford insight into the internal mechanisms underlying spike production; of particular importance here is the comparison of models of spike-production mechanisms, which is typically effected through comparison of the corresponding spike-train statistics. In interneuronal analysis, simultaneously recorded spike trains are compared statistically to reveal information concerning possible connections between neurons, shared sources of activity, responses to stimuli, and synaptic input-output relations. It is in the comparison of neuronal spike trains, however, that we believe the greatest promise of these techniques to lie, despite the greater incursion of mathematical difficulties when compared with single-train analysis.

In summarizing currently used techniques of spike-train analysis, we are further prompted by our feeling that these computational techniques have lacked adequate theoretical underpinning, which has resulted in (a) inconsistency of nomenclature and notation in this field, (b) difficulties, not always well enough appreciated, in assigning measures of statistical significance to experimental findings, (c) presentations of experimental data

in several forms that are in fact mathematically derivable from each other, and (d) the risk of attributing particular physiological significance to results that illustrate purely mathematical theorems or that are more plausibly attributable to chance effects. With the advent of the high-speed electronic digital computer, it has become feasible to perform on a routine basis the lengthy computations required for spike-train analysis; it is our feeling that in relating these computations to the underlying mathematical (i.e., probabilistic and statistical) theory, the potential usefulness of these techniques will be enhanced.

Basic Terminology and Scope

In every instance in which a detailed examination is made of the timing of neuronal events, we are forced to realize that a certain degree of unpredictability or randomness is present in the underlying process. For some purposes we can afford to overlook this aspect of the record, but for other purposes, and in particular those we focus on here, the question of interest, and indeed the source of greatest information about the process being observed, is the variability and randomness of the spike train.

This very property forces us to describe the spike train in statistical terms, and to view the processes underlying it either as inherently probabilistic or as

sufficiently complex that we can best and most simply treat them in probabilistic terms. Processes of this type are commonly referred to as stochastic processes. Indeed, the transfer and processing of information in nervous systems may be viewed as a repeated alternation, in time and space, of two different types of stochastic processes. The first type is characteristic of the continuous intra-neuronal fluctuations in significant state variables of each neuron. Typically, a state variable might be the membrane potential as observed at the primary spike-initiating locus, and a model of the underlying process might be described as a complicated type of random walk, with continuous time and "displacement" variables. The second type of stochastic process, which is the primary concern of this paper, arises in the study of the times of occurrence of interneuronally transmitted action potentials, i.e., the spike train, as commonly observed with either extracellular or intracellular microelectrodes. Because of our "all-or-none" conception of the nerve impulse, each spike is regarded as indistinguishable from the others produced by the same neuron. Furthermore, with each spike can be associated a unique instant of time, e.g., the time of maximum excursion of electrical potential, which can be measured with a high degree of precision. By virtue of the assumed indistinguishability* and instantaneity of the individual spike events, the stochastic process

*I.e., They are distinguishable only by where they occur in time.

characterizing the spike train can be considered an example of a stochastic point process. This process occurs in one dimension, corresponding to the time axis.

In any point process, in which all "events" (spikes, for example) are indistinguishable except for their times of occurrence, it is the elapsed times between events (e.g., the interspike intervals) that exhibit the properties of random variables. These intervals are regarded as being drawn (not necessarily independently) from an underlying probability distribution; if that distribution, together with its parameters, does not vary with time of observation, the stochastic point process is stationary.^{*} A sample from a stationary point process might be approximated, for example, by a spike train from a spontaneously firing neuron, or from a neuron well adapted to a steady stimulus. A monotonic trend in the firing rate or other parameter is one of the many features of a spike train that may preclude its characterization as stationary.

In most of this paper, spike trains will be considered as realizations of stationary point processes (except for local effects due to stimulation). We cover first the

^{*}Strictly speaking, stationarity is defined in terms of the invariance under translation in time of the joint distribution of numbers of events in fixed intervals of time (Cox & Miller, 1965, pp. 339-340; Cox & Lewis, 1966, p. 59). An equivalent definition in terms of interval distributions is difficult to formulate rigorously, mainly because of the complications introduced by the choice of the starting point for describing the process.

problems of description and interpretation as encountered with a single neuron exhibiting spontaneous or well-adapted activity. Next we consider the effects on a spike train of isolated, repeated presentations of a stimulus. Then we discuss the analysis of two simultaneously recorded spike trains, in the absence and in the presence of repetitive stimulation. Finally, we consider the effects of nonstationarity on the statistical measures described for stationary processes. The presentation is illustrated with examples drawn both from animal experiments and from digital-computer simulations (Perkel, 1965).

THE SINGLE SPIKE TRAIN

Stochastic Point Processes: Basic Nomenclature

A stochastic point process, as mentioned above, is a stochastic process "whose realizations consist of a series of point events" (Cox & Miller, 1965). The point events are considered to be instantaneous and indistinguishable (except for position in time); for neuronal spikes, therefore, we consider, for example, the time corresponding to the maximum of the observed action potential to be the time of occurrence, and we ignore all other characteristics of the spike, such as duration, amplitude, undershoot, etc.

In a stationary point process, the underlying probability distributions governing the times of occurrence of the point events do not vary with respect to an arbitrary translation of the time axis.* Therefore, accelerations and decelerations in firing rates, and effects such as fatigue and adaptation, disqualify spike trains from acceptance as realizations of stationary point processes. We consider the detection of nonstationarity and its effects, if present, in the final section of this paper. For a spike train observed in the absence of repetitive stimulation, the assumption of stationarity means, in a practical sense, a neuron that does not display any apparent trend in firing rate, and whose "mode"

* A more explicit definition is given by Cox & Lewis (1966), p. 59.

of firing does not exhibit any significant shift from one portion of the record to another.

One important class of stationary point processes, known as renewal processes, has the property that the lengths of intervals between events are statistically independent.* Neuronal spike trains rarely satisfy this requirement completely; even those spike trains that can adequately be described as stationary often exhibit serial dependence among interspike intervals.

Many results first established in renewal theory have subsequently been generalized to nonrenewal stationary point processes (McFadden, 1962), and in some cases even to nonstationary processes. The terminology of renewal theory, however, has been retained because of its intuitive appeal, and we use it here. Another set of metaphors more appropriate to neuron firings could easily be substituted.

The Poisson, the Erlang, and the Weibull processes are some of the most commonly encountered renewal processes with particularly simple properties. These are discussed amply in the literature (Cox, 1962) and have been applied to the description of certain classes of neuronal spike trains. Most spike trains with independent interval lengths,

* A renewal process is stationary only if observation begins at a random instant in time, in which case the interval from that instant to the first event has a different distribution from that of subsequent interevent intervals (Cox, 1962). We will assume this condition when discussing renewal processes.

however, do not fall into any of these mathematically attractive classes.

In the following sections we describe some of the most important and useful statistical measures of spike-train properties.

Order-independent Statistical Measures

For both renewal and nonrenewal stationary point processes, the (marginal) distribution of intervals between successive events is of paramount importance in characterizing the process. For a renewal process, in fact, the distribution of intervals completely characterizes the process. For finite samples of data, such as an observed neuronal spike train, the interspike-interval histogram serves as an estimator of the "actual" probability density function (pdf).

To construct it, the range of observed interval lengths is customarily divided into bins of equal width δ ; if the i^{th} observed interspike interval T_i satisfies the inequalities

$$(1) \quad (j - 1)\delta < T_i \leq j\delta,$$

then that interval is placed in bin j of the histogram.

The bins are thus numbered $1, 2, \dots, J$. Letting N_j designate the number of intervals placed in bin j in an observation of N intervals (i.e., $N + 1$ spikes), then the

ratios N_j/N are a smoothed estimate of the pdf $f(\tau)$; i.e., they estimate the corresponding integrals

$$(2) \quad N_j/N \approx \int_{(j-1)\delta}^{j\delta} f(\tau) d\tau.$$

This quantity is the probability that the duration of a randomly chosen interval lies between $(j - 1)\delta$ and $j\delta$. The estimator for the average value of the pdf within the bin is given by

$$(3) \quad f_j \equiv N_j/(N\delta).$$

Although N interval measurements are used in estimating the usual population parameters such as mean, variance, etc., it is only for a renewal process that the N observations are independent. Measures of precision assigned to these estimates by the standard formulas may be misleading if the process is not a renewal.

For both renewal and nonrenewal processes, there are several functions completely equivalent mathematically to the pdf $f(\tau)$ or its estimator f_j (Cox, 1962, pp. 2-7). These are illustrated in Fig. 1. One of these, the (cumulative) distribution function,

$$(4a) \quad F(\tau) = \int_0^{\tau} f(t) dt = \text{prob } (T \leq \tau),$$

is estimated by the empirical distribution function

$$(4b) \quad F_j = \sum_{k=1}^j f_k \delta$$

and, in neurophysiological terms, measures the probability that a neuron will have fired by time τ from the last firing.

Fig. 1

Equivalent form of the interspike-interval distribution. a-c, interspike interval histogram (estimate of the pdf). d-f, survivor function. g-i, hazard function. a, d, g drawn from a "pacemaker neuron": independently normally distributed intervals with mean 100 msec and standard deviation 5 msec. b, e, h from a Poisson process with mean interval 10 msec, conditioned by a normally distributed dead time (refractory period) with mean 10 msec, standard deviation 2 msec. c, f, i from a log-normal distribution, with pdf $x^{-1}(2\pi\alpha)^{-\frac{1}{2}} \exp \{-[\log(\rho x)]^2/(2\alpha)\}$, with $\rho = 27.2 \text{ sec}^{-1}$ and $\alpha = 2$. Each sample contains approximately 2000 intervals.




Fig. 1(a)




Fig. 1(b)




Fig. 1(c)





Fig. 1(d)






Fig. 1(f)




Fig. 1(g)




Fig. 1(h)




Fig. 1 (1)

The survivor function, the complement of $F(\tau)$, or

$$(5) \quad \mathcal{F}(\tau) = 1 - F(\tau) = \text{prob} (T > \tau),$$

is the probability that the neuron will not have fired by time τ .

A third function, the hazard function, measures (in the terminology of renewals) the instantaneous risk of failure of a component known to be of age τ . It is given by

$$(6) \quad \varphi(\tau) = f(\tau)/\mathcal{F}(\tau) = f(\tau)/[1 - F(\tau)].$$

In the neurophysiological context, the quantity $\varphi(\tau) \Delta\tau$ is the probability that a neuron will fire during the time interval $\Delta\tau$, given that it has not fired prior to the time τ . This function is called the "age-specific failure rate" by Cox (1962). In the neurophysiological literature, it is denoted as $c(x)$ by Poggio & Viernstein (1964), and called by them the "postimpulse probability"; the same function is called "conditional probability" by Goldberg et al. (1964). It is also denoted as $\lambda(x)$ by McGill (1963), who points out some earlier synonyms for the same function: "IRT/OPS," "conditional density function," and "hazard function."

For a Poisson process the hazard function is a constant (Fig. 1b, bottom). Other processes display positive or

negative "ageing" accordingly as $\varphi(\tau)$ increases or decreases with τ . Interspike-interval distributions from pacemaker neurons, for example, characteristically display positive ageing (i.e., an increased "hazard" of firing as a function of time since the last spike; Fig. 1a, bottom), whereas long-tailed distributions, such as those obtained for some neurons in the auditory system (Gerstein & Mandelbrot, 1964), display negative ageing (Fig. 1c, bottom). Note that the estimate of the hazard function loses precision for long intervals (Watson & Leadbetter, 1964a, 1964b).

Summaries of the interspike-interval distribution are furnished by various scalar quantities; estimation of these scalars from finite samples of spike-interval data does not differ in any essential respect from estimation using sample data from any other source. Useful quantities are the mean interval μ , the interval variance σ^2 , the standard deviation σ , and the coefficient of variation σ/μ . The mean firing rate ρ is defined as the reciprocal of the mean interval. Standard measures of skewness and kurtosis are often useful for describing and classifying interval distributions.

Order-dependent Statistical Measures

It is of considerable neurophysiological interest to determine whether or not successive interspike intervals are independent in the statistical sense, i.e., whether or not the spike train can be described as a realization of a

renewal process. Cox and Lewis (1966, pp. 164-171) discuss two classes of tests for independence of intervals, one based on sample serial correlation coefficients, and the other based on the spectrum of intervals. These authors state that "this is a difficult problem because the null hypothesis is very broad and the alternatives usually not at all clearly specified, and also because the associated distribution problems are hard." Moreover, if the intervals are not independent, such tests offer little information about the type of dependence, and they have so far found little application in spike-train analysis.

Aside from their use in hypothesis testing, certain statistical measures have been used to describe and quantify serial dependence among interspike intervals. We call these measure order-dependent, in order to distinguish them from those based solely on the interspike-interval distribution, discussed above. The methods of spectral analysis, recommended by Cox and Lewis (1966; see also Bartlett, 1963), have not been used for analyzing spike-interval data, but their potential utility deserves extensive investigation. Serial correlation coefficients are statistics based on joint distributions of intervals, and another set of measures is based on intervals between nonadjacent spikes. We describe here the latter two classes of techniques.

Use of the joint interval density for spike-train analysis was introduced by Rodieck, Kiang & Gerstein (1962). The data are displayed in the form of a scatter diagram, in which the length of an

interspike interval is represented by the abscissa, and the length of the next interval in the record by the ordinate. Each point on the diagram then represents a pair of adjacent intervals. An alternate form is the matrix equivalent, Smith & Smith (1965). If successive intervals are independently distributed, then the normalized frequency distribution along the ordinate is the same for each abscissa value, and vice versa. This implies that the corresponding row and column means in the corresponding joint interval histogram have a constant expected value. An observed constancy of row and column means, which is a necessary condition for independence of adjacent intervals, has in practice been used as a sufficient test. Departures from independence are reflected not only in these means but also in the symmetries of the scatter diagram itself. For example, an overall "upward" trend in the joint-interval scatter diagram (as illustrated in Fig. 7 of Rodieck et al. (1962), units 261-1, R-4-10, and 240-1) indicates positive correlations between successive intervals. This means, loosely, that short intervals tend to be followed by short ones, and long intervals by long ones. A "downward" trend would imply negative serial correlation. The uncorrelated case is illustrated by unit 259-2, in the same figure. The corresponding row and column means are shown in Fig. 8 of the same work.

A quantitative measure of such correlation is furnished by the serial correlation coefficient of interval lengths,

which is defined as follows: We define the covariance of interval lengths, of lag j , by

$$(7) \quad C_j = E[(T_i - \mu)(T_{i+j} - \mu)] (j = \dots, -1, 0, 1, \dots),$$

where T_i is the i^{th} interspike interval in an (infinite) stationary train of spikes, with mean interval μ and variance σ^2 . Then the serial correlation coefficient ρ_j of order j is the ratio of the corresponding covariance to the interval variance:

$$(8) \quad \rho_j = C_j / \sigma^2.$$

In a finite sample, the mean μ and the variance σ^2 must be estimated from the sample data. To avoid the slight bias introduced by use of sample mean and variance to estimate the corresponding population parameters, more complicated formulas are suggested by Cox and Lewis (1966, pp. 89-92).

The serial correlation coefficient of lag 1 furnishes a single scalar quantity as a summary of the entire diagram of the joint interval distribution of lag 1. Joint interval distributions of larger lags (i.e., for nonadjacent intervals) have not been used in the analysis of spike trains; the corresponding serial correlation coefficients, however, have been used extensively. The set of serial correlation coefficients is

usually called the serial correlogram; it has sometimes (e.g., in Hagiwara, 1954) been called the "autocorrelogram" or "autocorrelation," terms that we reserve for a different function (see below).

The expected value of the serial correlation coefficients of all orders (lags) is approximately zero if the intervals arise from a (stationary) renewal process, i.e., if the intervals are drawn independently from a common distribution. For large N , for a renewal process, the quantity $\rho_i/(N - 1)^{1/2}$ has a unit normal distribution. However, the distribution of the sample serial correlation coefficient is not known for small samples, and the sample coefficients of various lags are correlated for moderate sized samples (Cox & Lewis, 1966 p. 165). No test of independence of intervals is known which is based jointly on several serial correlation coefficients (P. A. W. Lewis, private communication).

One useful expedient, however, is to subject the sample of interspike intervals to random shuffling, which destroys serial dependence but preserves the order-independent statistics of the sample. Shuffling thus converts the sample to one from the corresponding renewal process. The recomputed serial correlogram for the shuffled train provides a control case, in which the departures from zero of the serial correlation coefficients are in fact those due to random fluctuations. The net departure may be measured, for example, by the sum of squares of the coefficients. In principle, the shuffling and recomputation can be repeated at length, to

provide an empirical sampling distribution of the sum of squares under the null hypothesis of serial independence. From this distribution, tests of the independence hypothesis for the unshuffled data can readily be constructed. This kind of procedure is discussed by Cox & Lewis (1966, p. 165) as a permutation test of serial correlation. A refinement discussed there is to replace the observed interval values by ranks or exponential scores. In this way the sampling distribution can be computed once and for all for a given sample size.

The most frequently encountered source of positive contributions to the serial correlation coefficients is a long-term trend in the data; a sufficiently great monotonic increase or decrease in the firing rate over the time of observation will contribute a positive component to each serial correlation coefficient, out to lags of arbitrarily high order. Monotonic trends, of course, are a form of nonstationarity; they are discussed further in a subsequent section.

In a stationary point process that is not a renewal process, the serial correlogram furnishes indications as to the nature of departure from independence among intervals. Local trends in firing rate will introduce positive contributions to the serial correlation coefficients, primarily to the lower lag values (Hagiwara, 1949; Junge & Moore, 1966). Cyclic variations in firing rate, for example,

produce a damped oscillation in the serial correlogram, which starts at positive values. Somewhat similar oscillations are produced by fairly regular occurring bursts of spikes, as exemplified by medullary respiratory neurons. Irregular bursts of spikes, such as have been described by Smith & Smith (1965), are characterized by positive serial correlations for low lags, followed by slightly negative and then zero correlation coefficients. Alternation between long and short intervals, as commonly seen in certain cells in the dorsal column nuclei (Amassian et al., 1962) or in locust wing-muscle motor neurons (Wilson, 1964), gives a strongly negative first serial correlation coefficient with subsequent alternation in sign of the coefficients of higher order. Other types of patterned activity have corresponding signatures in the serial correlogram. Some examples, from computer simulations, are illustrated in Fig. 2.

It must be emphasized that both the joint interval histogram and the serial correlogram must be interpreted with caution, and in conjunction with other statistical measures. The positive contribution to low-order correlation coefficients due to local or global trends in firing rate may mask any negative correlation between adjacent intervals that would otherwise be apparent. Segmentation of the data may clarify this situation, but at the expense, of course, of statistical reliability. A gap in the record, which

Fig. 2.

Typical serial correlograms. a. Independent intervals drawn from a Weibull distribution, with pdf $\alpha \rho (\rho x)^{\alpha-1} \exp -(\rho x)^\alpha$; $\rho = 15 \text{ sec}^{-1}$, $\alpha = 0.6$, 2161 spikes. b. Simulated neuron producing irregular bursts. Resting level of membrane potential -70 mv , reset to -100 mv after a spike; exponential recovery with decay constant of 6.93 sec^{-1} . Asymptotic threshold -40 mv , reset to -20 mv after a spike; exponential recovery with decay constant 3.47 sec^{-1} . Two input channels producing 18.5-mv EPSP's: one channel with mean firing rate $2.0/\text{sec}$, the other with mean firing rate $2.1/\text{sec}$; standard deviation of intervals 3% of mean. Cell fires with mean interval 2.6 sec , standard deviation 2.9 sec , with sample of 389 spikes. c. Simulated neuron producing more regular bursts. Similar to case b, but with standard deviation of intervals in input channels 0.6% of respective means. Cell fires in bursts of four or five spikes; interspike intervals within bursts 0.5 to 1.0 sec ; bursts start at intervals of approximately 10 sec . Sample of 424 spikes has mean interval 2.4 sec , standard deviation 2.8 sec . d. Decelerating train. Intervals drawn from a time-dependent normal distribution with $\mu(t) = 0.1 - 0.02 e^{-.01t} \text{ sec}$, $\sigma = 0.01 \text{ sec}$. Sample of 2447 spikes. e. Alternation between long and short interval lengths. Intervals drawn alternately from distribution with mean 0.11 sec and from distribution with mean 0.09 sec , each with standard deviation 0.015 sec . 5000 spikes.




Fig. 2(a)




Fig. 2(b)




Fig. 2(c)




Fig. 2(d)




Fig. 2(e)

introduces an exceptionally long interval, may grossly affect the interval variance and seriously distort the serial correlogram; this effect is most pronounced in pacemaker neurons (small coefficient of variation of intervals). On the other hand, effects of trends and gaps on the serial correlogram are much less pronounced in data with a large inherent variability (large coefficient of variation), such as a Poisson process with a time-varying rate parameter, or neurons with highly irregular firing times. It is important to measure and correct for these distorting effects since significant information about the physiology of the neuron--such as refractory effects, persistence of synaptic effects, etc.--may be uncovered through correlational analysis of successive intervals (Firth, 1966; Junge & Moore, 1966; Geisler & Goldberg, 1966).

The use of ranks or exponential scores greatly alleviates the effects of gaps in the record. The use of the estimated spectrum of intervals overcomes some of the effects of trends. The tests for independence of intervals based upon the spectrum of intervals, as described by Cox & Lewis (1966, pp. 67 et seq.), appear to have been neglected by investigators of spike trains. The estimated spectrum gives a single test for independence based on all of the data, and thereby overcomes the obstacles mentioned above to the interpretation of the serial correlogram. Computation of the spectrum has been greatly facilitated by a recently

devised algorithm (Cooley & Tukey, 1965), which has been incorporated in a set of computer programs by P. A. W. Lewis (Lewis, 1966) for the statistical analysis of series of events.

Joint interval distributions and the corresponding serial correlation coefficients involve time intervals that are defined by two successive spikes. The second class of order-dependent statistical measures that we discuss involves time intervals between nonsuccessive events. Denoting as a first-order interval the elapsed time from an event to the next following event, we may define a second-order interval as the elapsed time between an event and the second following event, etc. An n^{th} -order interval is the sum of n consecutive first-order intervals, and is spanned by $(n + 1)$ consecutive spikes (see Fig. 3).

The probability density of the n^{th} -order interval is designated $f_n(\tau)$. The interval density for successive events is thus the first-order density: $f_1(\tau) = f(\tau)$.

In the special case of a renewal process, the higher-order densities may be obtained by successive convolutions of the first-order density because the intervals are independent. Thus, the second-order interval density is given by the convolution integral

$$(9a) \quad f_2(\tau) = \int_0^{\infty} f(t)f(\tau - t) dt,$$

Fig. 3.

Higher-order interspike intervals. A first-order interval is the time difference between adjacent spikes. A second-order interval lies between a spike and the second spike following, etc. Note that an interval of order n spans $n + 1$ spikes. See text.




Fig. 3

and in general we have the recursion

$$(9b) \quad f_{n+1}(\tau) = \int_0^{\infty} f_n(t)f(\tau - t) dt.$$

These equations do not apply when successive interval durations are not independent.

A related function is the renewal density, $h(\tau)$, which specifies the probability of encountering any event as a function of time after a given event; i.e.,

$$(10) \quad h(\tau) = \lim_{\Delta\tau \rightarrow 0} \text{prob} \left\{ \begin{array}{l} \text{an event in } (\tau, \tau + \Delta\tau) \\ | \text{ an event at } 0 \end{array} \right\} / \Delta\tau.$$

Since any event encountered must be either the first, second, ..., etc., event after the event at time 0, it is evident that the renewal density is the sum of the interval densities of all orders:

$$(11) \quad h(\tau) = \sum_{k=1}^{\infty} f_k(\tau).$$

This is illustrated in Fig. 4, in which is shown the renewal density (as estimated by a histogram) from a sample of spike activity of a computer-simulated neuron; together with it are displayed the corresponding interval densities of the first four orders.

In spike-train analysis, the renewal density is often called the autocorrelation (Gerstein & Kiang, 1960), since

Fig. 4.

Composition of the autocorrelation (renewal density). Computer simulated noisy pacemaker with inhibitory synaptic input. Mean interspike interval without inhibition 0.47 sec; Poisson arrivals of IPSP's at mean rate of 4/sec; maximum hyperpolarization of IPSP's normally distributed with mean 5 mv, standard deviation 1 mv. Mean interspike interval 0.54 sec, standard deviation 0.09 sec, 2000 spikes.

a. Autocorrelation histogram. b. Interspike-interval histogram. c-f. Second- through fifth-order interval histograms. See text.




Fig. 4(a)




Fig. 4(b)




Fig. 4(c)




Fig. 4(d)




Fig. 4(e)




Fig. 4(f)

if the spike train is regarded as a signal of zero amplitude everywhere except where a spike is present, and if each spike is represented as a Dirac delta function, then the renewal density corresponds to the autocorrelation as ordinarily defined for continuous signals. In the usage of Cox & Lewis (1966), the renewal density is one of a class of "intensity functions." Another synonym that has gained some currency is "expectation density"; it is used by Poggio & Viernstein (1964) and others, following a usage introduced by Huggins (1957). Yet another synonym is "post-firing interval distribution," introduced by Lamarre & Raynauld (1965).

Another property of the autocorrelation is that it "flattens out" to a constant value; i.e.,

$$(12) \quad \lim_{\tau \rightarrow \infty} h(\tau) = 1/\mu.$$

In general, this limit is approached slowly for pacemaker neurons (narrow interspike-interval densities), and more rapidly for spike trains with greater variation in interval length. The limiting value is attained for all values of τ for a Poisson process; it can be shown, in fact, that a Poisson process is the (ordinary) renewal process having a constant renewal density (Cox, 1962).

The representation of higher-order interval densities as corresponding convolutions of the first-order density (Eq. 9) holds only for independent intervals. If there is

serial dependence (as measured by the serial correlogram); more complicated expressions are required, as given by McFadden (1962). Thus the sum of the interval densities of various orders (the autocorrelation) for two spike-train sequences having identical interval distributions will be different if in one train the intervals are independent (i.e., a renewal process) and in the other they are dependent. Hence it is possible to compare the observed autocorrelation with that predicted under the independence hypothesis as a test of that hypothesis. One convenient computer method is that of prolonged random shuffling of the intervals, as discussed above. The autocorrelation of the shuffled train then represents a control case of serial independence. Discrepancies between the unshuffled and shuffled autocorrelation not only furnish a test of serial dependence, but may also indicate the nature of that dependence.

Shuffling may either enhance or flatten peaks in the autocorrelation. For example, if interspike-interval lengths exhibit negative serial correlation, peaks in the autocorrelation are generally broadened by shuffling. On the other hand, if interval lengths are positively correlated, shuffling of intervals may sharpen peaks in the autocorrelation. An example of the latter effect is shown in Fig. 5, in which the original interval sequence was generated by a three-stage semi-Markov process (Cox & Lewis, 1966, p. 82), according to which relatively long intervals were more likely to be followed by long intervals than by short ones, and vice versa.

Fig. 5.

Effect of shuffling on peaks in the autocorrelation. Intervals generated by a semi-Markov process with three states with mean intervals as follows: state 1 (short), .09 sec; state 2 (medium), .10 sec; state 3 (long), .11 sec. All intervals normally distributed with standard deviation .01 sec. Transition matrix

$$\begin{pmatrix} 0.70 & 0.20 & 0.10 \\ 0.45 & 0.10 & 0.45 \\ 0.10 & 0.20 & 0.70 \end{pmatrix}$$

is such that a long interval is most likely to be followed by a long interval, and a short by a short, Mean interval 0.10 sec, standard deviation 0.014 sec, in sample of 2000 spikes; first five serial correlation coefficients: 0.251, 0.184, 0.085, 0.071, 0.040. After shuffling of intervals, first five serial correlation coefficients were -0.018, 0.017, 0.021, -0.041, -0.004. a. Autocorrelation histogram, unshuffled data. b. Autocorrelation histogram of reconstructed spike train after shuffling of interspike intervals. Note sharpening of peaks. See text.




Fig. 5(a)




Fig. 5(b)

To avoid confusion it should be emphasized that the autocorrelation is a function of time; the serial correlogram (Hagiwara's (1954) "autocorrelation") is a function of the serial position of the interval, an integer. The two functions need not correspond at all. For example, a pacemaker cell that fires at nearly uniform intervals will have a strongly oscillating autocorrelation, whereas the serial correlogram may be positive, negative, oscillatory, or zero.

Long-term trends or slow oscillations in firing rates are ordinarily not obvious in the autocorrelation, but are more typically revealed in the serial correlogram. These and other effects of relatively long-term rate variations are discussed below, in the section dealing with the effects of nonstationarity.

Description, Prediction, and Information

Having presented certain statistical measures that can be applied to individual spike trains, we now introduce briefly some considerations about the adequacy and utility of statistical descriptions of the train. These considerations bear on the use of statistical descriptions both in characterizing and classifying neurons, and in comparing observed spike data with those predicted by models.

The simplicity of statistical description of a spike train differs widely from case to case. A single parameter suffices

to describe a Poisson process, whereas two are required for an Erlang or Weibull process. If a spike train fits into one of these categories, not only is the characterization of the particular spike train extremely simple, but also the pdf of the interval distribution (and therefore the hazard function, autocorrelation, etc.) may be written as an explicit equation. For less easily described renewal processes, the entire interspike-interval histogram is required to characterize its properties. When successive intervals are not independent, then much more complicated descriptions are needed, except in special cases of highly patterned spike configurations. When higher-order joint interval densities or similarly elaborate statistical measures are required to effect a reasonably complete description, the statistical description itself is unmanageable, and its practical utility is highly questionable. The limit is reached, of course, when the number of parameters in the statistical model equals (or exceeds!) the number of spikes in the sample, at which point statistical analysis loses all justification.

The relationship between the statistical properties of a spike train and the information-handling capability of the neuron is generally complex. Estimates of channel capacity depend strongly upon the particular choice of encoding scheme inputted to the neural structure. Once a choice of encoding scheme has been made, then estimates of

channel capacity can be obtained on the basis of interval statistics. For example, if a parameter is thought to be encoded in terms of mean firing rate, then a smaller coefficient of variation of intervals gives rise to a higher channel capacity, etc. Interval statistics, however, cannot of themselves provide a choice of coding scheme. Facile "derivations" of information-handling characteristics of neurons based solely or primarily on spike-interval statistics are usually misleading or worse. For a more extended discussion of these problems, the reader is referred to Moore et al. (1966) and to Segundo et al. (1966).

SINGLE SPIKE TRAINS IN THE PRESENCE OF STIMULATION

In many neurophysiological experiments a controlled series of changes in the physical environment is introduced. We consider here the case of a repeated, relatively short stimulus. In order to detect and evaluate the effect of such a stimulus train on the train of spikes, it is now common to compute a post-stimulus-time (PST) histogram (Gerstein & Kiang, 1960). Specifically, the PST histogram shows the probability of firing as a function of time after the stimulus onset. As shown below, this measure is equivalent to a cross correlation between the train of stimulus presentations and the train of spikes. If the stimulus has no effect on the pattern of the spike train, the PST histogram will be flat (subject to the usual statistical fluctuations). On the other hand, if the stimulus does produce a time-locked "evoked response" in the spike-train pattern, the PST histogram will show deviations from flatness. A peak in a PST histogram indicates a higher probability of firing at that particular time after stimulation and can presumably be associated with an excitatory process. Dips in a PST histogram indicate a lower time-locked probability of firing and often are associated with inhibitory or refractory processes.

In interpreting a PST histogram it is obviously necessary to decide on the statistical significance of the observed deviations from flatness. A simple method that

has been used (Weiss, 1964) is to compute the mean square deviation of all bins from the mean level of the histogram. Some criterion value for this number can be chosen to distinguish a "flat" PST histogram from one that shows a weak time-locked response pattern.

There is some difficulty in this type of significance test, however, since successive bins in the histogram may not represent independent quantities. For example, since each firing of the neuron is followed by a refractory period, there is always a tendency toward a negative correlation between adjacent bins in the histogram. An empirical control case can be constructed by randomly shuffling the intervals of the spike train and computing the corresponding "flat" PST histogram. The distribution of mean square deviations from mean bin level can be obtained from a set of replications of this procedure, and the criterion value can then be chosen as usual to satisfy a specified error probability. A control case can also be constructed using fictitious times of stimulus presentation in a portion of record where no actual stimulations were presented (Gerstein, 1960; Burns & Smith, 1962).

Because of such correlation effects, it is necessary to verify that features in a PST (or similar) histogram that are suspected of having significance are not simply artifacts of the choice of bin width. Two techniques that may be helpful are (1) recomputing the histogram with a different bin

width that is not a simple fraction of the original bin width, and (2) calculating the autocorrelation of the PST histogram. Meaningful features should have a width of several bins. The existence of "wide" features in the PST histogram will be shown by large values near the origin of its autocorrelation function.

Another way of analyzing stimulus effects on a single spike train is to measure the elapsed times or "latencies" between stimulus presentations and the earliest encountered subsequent spikes. Latency and PST studies are discussed in the review paper by Moore et al. (1966).

TWO SIMULTANEOUS SPIKE TRAINS

The Problem of Functional Relationship

It has become increasingly feasible to record spike trains simultaneously from several neurons (Gerstein and Clark, 1964; Simon, 1965). In the interpretation of these simultaneous spike-train recordings, the trains may be compared by pairs.* The first question to be asked is whether the two trains of spikes are independent. More specifically, we wish to make a statistical test of the null hypothesis that the two trains are drawn from independent point processes. This would imply that the two neurons are functionally unrelated. The test is accomplished, in principle, by computing a histogram, which estimates a suitable function (e.g., either the cross density or the cross correlation function, described below), and comparing the estimate with the predicted function as based on the assumed independence of the two trains. If the observed and predicted functions differ significantly, the null hypothesis is rejected, and the trains are considered to be dependent.

An observed dependence between two spike trains can arise from one (or both) of two sources: (1) functional interaction and (2) common input. By functional interaction, we mean any mechanism by which the firing of one neuron influences the probability of firing by the other neuron. Such mechanisms could be synaptic (whether direct or mediated through interneurons), ephaptic, or due to "field effects."

*Statistical measures for the concurrent intercomparison of three or more trains are currently under investigation.

By common input we mean any mechanism that simultaneously modulates the firing patterns of both neurons. Such mechanisms could involve synaptic contact from branches of the same axon, or field effects from a source other than the two neurons. It must be emphasized that long-term concomitant changes in firing rates, if shared wholly or in part by two neurons, constitute a form of dependence, and, if sufficiently pronounced, will be detected by statistical test procedures (see discussion below, on Types of Dependence).

Independence: Consequences of the Null Hypothesis

According to the null hypothesis, spike trains A and B are independent in the mathematical sense. This means that spikes in train A occur at moments taken at random with respect to train B. In relating the two spike trains we may therefore use some mathematical results about single point processes observed from random moments in time.

One such result concerns the distribution of so-called recurrence times, which are defined as follows. From a random instant in time, we denote by V_1 the time until the next event in a point process, and by V_{-1} the time backward to the most recent event in the point process; V_1 and V_{-1} are known as the forward and backward recurrence times, respectively (Fig. 6). Each of these times has the same distribution with the pdf

$$(13) \quad g_1(\tau) = [1 - F(\tau)]/\mu = \mathcal{F}(\tau)/\mu,$$

Fig. 6.

Recurrence and waiting times. Train B is a train of spikes. If the events in train A are random instances in time, then V_1 and V_{-1} are the (first-order) forward and backward recurrence times, respectively; V_2 and V_{-2} are the corresponding second-order recurrence times, etc. If train A is a spike train, then (in the usage of this study) the corresponding time intervals are called waiting times, and are designated W_1 , W_{-1} , W_2 , etc. See text.




Fig. 6

where μ is the mean interspike interval; $F(\tau)$, as before, is the cumulative probability distribution of the intervals, and $\bar{F}(\tau)$ is the survivor function. This result is true for both renewal and nonrenewal processes (Cox, 1962; McFadden, 1962).

In terms of two spike trains, this result means that the time from a randomly selected spike in A to the immediately following, or to the immediately preceding, spike in B is distributed with the pdf

$$(14) \quad \eta_1(\tau) = \eta_{-1}(\tau) = [1 - F_B(\tau)]/\mu_B$$

if the two spike trains are independent. A similar relationship holds, of course, for the times between any spike in B and the neighboring spikes in A, with the appropriate change of subscripts.

Generalizing this notion, we may make the corresponding measurements from two simultaneous spike trains, whether or not they are in fact independent. We designate as W_1 the "waiting time" from a spike in train A to the next subsequent spike in train B, and as W_{-1} the time backward to the most recent spike in train B. The distributions of these random variables are specified by their pdf's, $\eta_1(\tau)$ and $\eta_{-1}(\tau)$, which we call the forward and backward cross-interval densities, respectively. These densities, as usual, are estimated by histograms constructed from the observed spike trains. If

the trains are independent, the cross-interval histograms will agree with their predicted forms, according to Eq. 14.

It is to be noted that this prediction is based only on the interval distribution in train B. If the observed histograms agree with the predictions, however, it can be concluded that the trains are independent only insofar as adjacent spikes are concerned. Effects of one neuron's spikes on the other's spikes may well be delayed to the extent that they will not be revealed in the cross-interval histograms.

Departures of the cross-interval histograms from their predicted form may often indicate the type of dependence between the two trains. Furthermore, in cases of indicated dependence, the forward and backward cross-interval histograms will in general be different.

The second mathematical result deals not only with the first-order recurrence times (forward and backward to the adjacent events), but with the sum of these for all orders.

Letting V_2, V_3, \dots refer to the times from a random moment to the second, third, \dots events encountered in the process B, we may define in general $g_i(\tau)$ to be the pdf of recurrence time V_i . If these pdf's exist (as they do for most physically realizable processes, including spike trains), then it follows from a result of McFadden (1962, Eq. 2.21) that the sum of the recurrence densities of all orders is a constant:

$$(15) \quad G_B(\tau) \equiv \sum_{i=1}^{\infty} g_i(\tau) = 1/\mu_B.$$

Here μ_B is the mean interval between events in the process B, which is not necessarily a renewal process. Since the backward recurrence times V_{-i} of B have the same distributions as the corresponding forward recurrence times, their sum also has the same distribution as that for the forward times.

We proceed as before to generalize to the case of not necessarily independent spike trains, and we define the forward and backward waiting times of order i , W_i and W_{-i} , as the time measured from a spike in train A to the i^{th} subsequent spike encountered in train B, or backward to the i^{th} previous spike in train B, respectively (Fig. 6). The corresponding cross-interval densities of order i are denoted by $\eta_i(\tau)$. We further define the cross-correlation function $\zeta_{AB}(\tau)$ as the sum of all orders of cross-interval densities:

$$(16) \quad \zeta_{AB}(\tau) = \sum_{\substack{i=-\infty \\ i \neq 0 \\ i=\infty}} \eta_i(\tau).$$

Now, reasoning as we did with Eqs. 10 and 11, for the autocorrelation, we realize that for a given spike in train A, the time, backward or forward, to any spike in train B must be the waiting time of some order i ($i = \dots, -2, -1, 1, 2, \dots$). Therefore, the cross correlation represents the probability of encountering any event in train B as a function of time before or after an actual event in train A:

$$(17) \quad \zeta_{AB}(\tau) = \lim_{\Delta t \rightarrow 0} \text{prob} \{ \text{an event in B in } (t_0 + \tau, t_0 + \tau + \Delta t) | \\ \text{an event in A at } t_0 \} / \Delta t.$$

Functions of this kind are called "cross intensity functions" by Cox & Lewis (1966, p. 247).

The same reasoning applies, of course, if train A consists of random instants of time, and so we may write

$$(18) \quad G_B(\tau) = \lim_{\Delta t \rightarrow 0} \text{prob} \{ \text{an event in B in } (t_0 + \tau, t_0 + \tau + \Delta t) | \\ \text{random instant } t_0 \} / \Delta t \\ = 1/\mu_B.$$

There is a parallelism between the function $G_B(\tau)$ and the renewal density (autocorrelation) $h(\tau)$. If we observe a stochastic point process starting at a random event, then the probability of encountering spikes at any time thereafter is measured by $h(\tau)$. For large τ (i.e., at long times after the initial observation point), this density becomes a constant, equal to the mean firing rate $1/\mu$ (Eq. 12). On the other hand, if we start our observation at a random instant in time, rather than at a particular event, the probability density $G_B(\tau)$ of encountering a spike at some later time is already a constant. Heuristically, this means that if we know that a spike has occurred, we can predict (from the statistical properties of the spike

train) when subsequent spikes will occur; our prediction becomes worse the farther removed we are from the initially observed spike, until our prediction is based solely on the mean firing rate. But if we start at a random moment, we do not even know the position of one spike, and we can do no better at predicting subsequent spike locations than by using the mean rate. These considerations apply to both renewal and nonrenewal point processes.

The cross-correlation function $c_{AB}(\tau)$ has a useful symmetry relationship. Consider the compound event "a spike in record A is followed at a time X by a spike in record B"; this event may be equivalently described as the event "a spike in record B is preceded at a time X by a spike in record A." The probability of this compound event may be expressed in either of two ways, implying the following symmetry relationship for the cross correlation:

$$(19) \quad c_{AB}(\tau)/\mu_A = c_{BA}(-\tau)/\mu_B.$$

This identity holds even if the two trains are not independent. For this reason, the cross-correlation function is customarily measured in only one direction, from train A to train B, and for both forward and backward waiting times—i.e., for both positive and negative values of τ .

Application to Experimental Data

Thus, in order to test whether two spike trains are independent, we may make the following measurements:

(1) The cross-interval histogram. We select spikes in train A and construct a histogram of the times to the nearest spikes in train B. This histogram can then be compared with the estimated backward and forward recurrence-time densities $\eta_1(\tau)$, and $\eta_{-1}(\tau)$, which in turn are estimated from the interspike-interval histogram of train B, using Eqs. 4b and 14. This procedure must be repeated, after interchanging the rôles of the trains.

(2) The cross-correlation histogram. We select spikes in train A and construct a histogram of the times to all spikes in train B, both forward and backward, out to some specified time. This histogram provides an estimate of the cross-correlation function $\zeta_{AB}(\tau)$, which can be compared with the predicted constant value $1/\mu_B$.

If the two spike trains are independent, then the forward and backward cross-interval histograms will be equal to the prediction $[1 - F_B(\tau)]/\mu_B$, and hence equal to each other, within sampling effects. Since these equalities are a necessary but not sufficient condition for independence of the spike trains, the cross-interval histograms have two major uses: (1) as a corroboration of independence when indicated by a flat cross correlation, and (2) as a means of exploring suspected short-latency interactions, i.e., those occurring with a latency smaller than the mean interspike interval in train B. Cross-interval histograms

must be computed separately for train A against train B and vice versa. These measurements are related to the phase histograms of Wyman (1965).

Satisfactory statistical tests of the null hypothesis of independent firings have not yet been completely developed.

(See, for example, Cox & Lewis (1966), pp. 247-248.) There

are two sources of difficulty. The first arises from the lack of independence of successive bins in the cross-correlation histogram, as has been mentioned for the PST histogram. An observation contributing to one bin of the histogram is produced by the firing of a neuron; the refractoriness of the neuron will then lower the probability of a contribution to the next bin. The magnitude of this effect depends on the degree and duration of refractoriness, as well as on bin width. If two neurons are in fact independent, consider a bin that has a larger number of spikes than average due to a random fluctuation. The next bin is then more likely to have a smaller number of spikes than average. Thus, the contribution to the sum of squares of deviations from the mean will be larger than if bins represented independent observations. This can lead to false attributions of dependence to cells that are, in fact, firing independently.

A related source of the same type of error arises from the sampling requirements for spike trains of neurons that are pacemakers. Two independently firing pacemakers of nearly the same frequency typically exhibit peaks in the

cross-correlation histogram, which oscillate with the period of the pacemakers. The magnitudes and phases of these peaks vary randomly from sample to sample, and they may disappear only with extraordinarily long samples. In Fig. 7, we show examples of cross-correlation histograms taken from two simulated independent pacemakers (normally distributed intervals with standard deviation of interval set at 10 percent of the mean). Each sample consisted of 4000 spikes in each neuron. The examples show the appearance and occasional disappearance of the peaks, which may be interpreted as a run of phase "locking" arising by chance. It is consequently very difficult to decide from reasonably sized samples whether two pacemakers are independent.

Thus, there are appreciable risks of falsely attributing dependence to independent spike trains, particularly when both neurons are pacemakers. The other possible type of error is false attribution of independence to trains that are in fact dependent. This is much less likely and can arise in either of two ways: (1) The dependence may be so weak that its effects are indistinguishable from "noise" in the sample taken. (2) A simultaneous combination of positive and negative interactions may coincidentally combine in such a fashion as to cancel each other. If this should happen, however, the two neurons may still be regarded as firing independently since, in an operational sense, the firing of one cannot be used to predict firing times of the others.

Fig. 7.

Cross-correlation histograms for two independent pacemakers. Each train has independently drawn normally distributed intervals, with mean 1.0 sec, standard deviation 0.1 sec; each train has 4000 spikes. a. Unshuffled data. b. After prolonged random shuffling of intervals in train B. c. After further prolonged random shuffling of intervals in train B. See text.


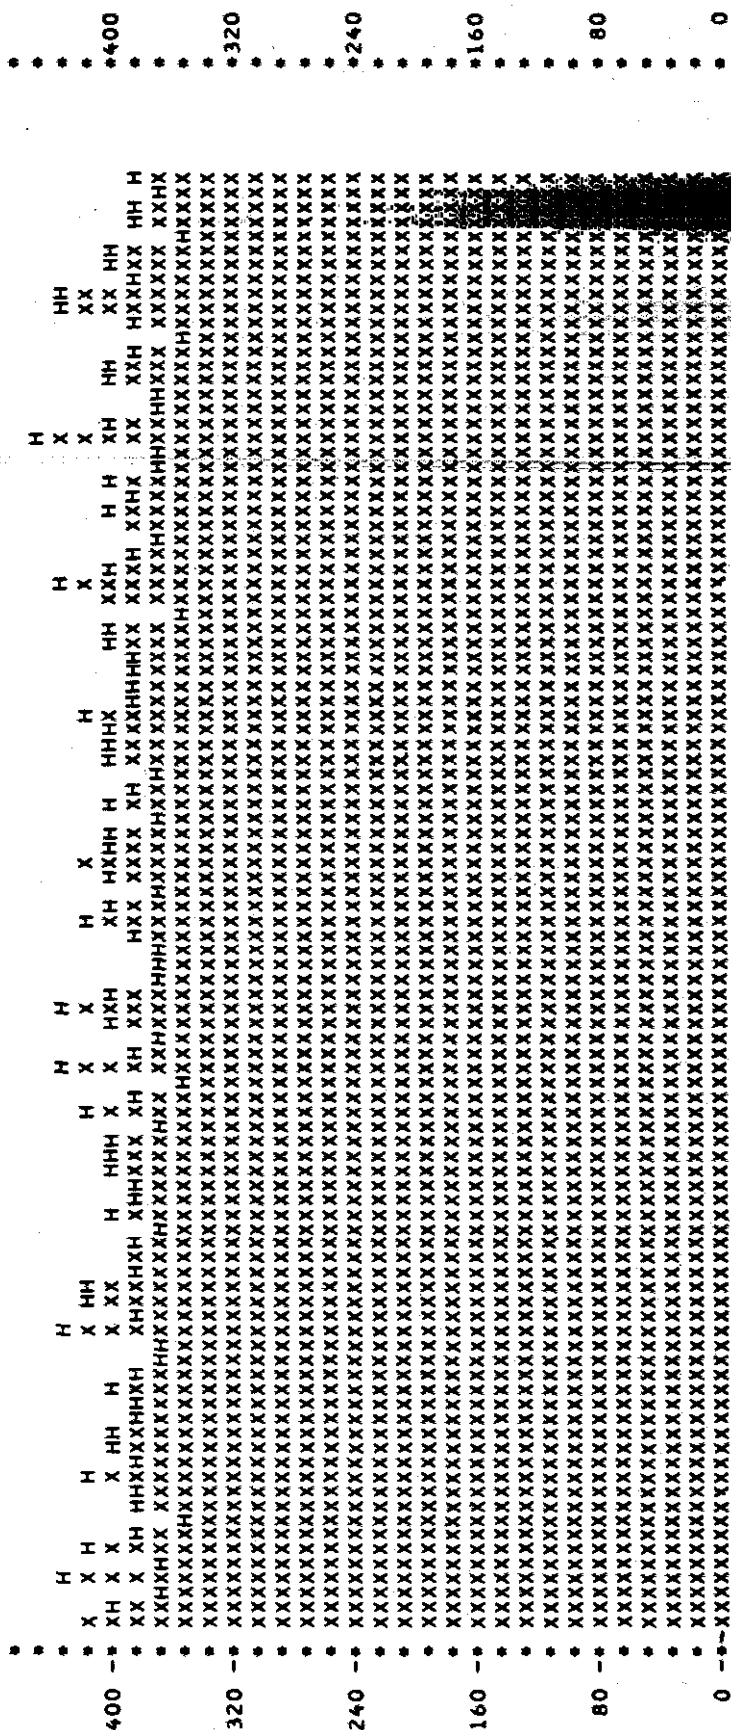


Fig. 7 (a)



0

0

-5.0

Lag (sec)

Fig. 7(b)




Fig. 7(c)

Types of Dependence

Dependence may arise from a number of physiologically important processes. We have already distinguished interaction and common input. The combinations of possibilities are virtually endless:

Two cells may interact through synaptic connections that may be either direct (e.g., an axon collateral of one of the cells forming a synapse with the other) or indirect (i.e., mediated through one or more interneurons). The effects of such interactions on the cross correlation cannot be predicted precisely without a detailed knowledge of the intracellular processes, including postsynaptic potentials, which underlie the production of spikes by each cell. In simple cases, however, the gross effects are intuitively obvious. For example, if cell A makes a single excitatory synaptic connection with cell B, with a mean conduction time w , then we would expect the probability that cell B fires to be enhanced during a period starting at a time w after the occurrence of every spike in cell A. Therefore, we observe a peak in the cross-correlation function $c_{AB}(\tau)$ near the point $\tau = w$. The shape of the peak depends on the details of the synaptic interaction (Fig. 8a). If the connection is inhibitory, a depression rather than a peak will be observed (Fig. 8b).

Two cells may receive input from a common source either directly or indirectly. In a simple example, the

Fig. 8.

Cross-correlation histograms for interacting neurons; data from computer simulations. a. Direct excitation. Neurons with characteristics similar to those of Fig. 2b. Each neuron excited by an independent Poisson source of EPSP's, amplitude normally distributed with mean 7.5 mv, standard deviation 1.0 mv; mean arrival rate 100/sec. Cell A makes excitatory synapse on cell B with latency of 100 msec; EPSP mean amplitude 10 mv, standard deviation 1 mv. In 250 sec, 2053 spikes produced by cell A, 2286 spikes by cell B. b. Direct inhibition. Same as Fig. 8a except A-B synapse is inhibitory (mean amplitude -10 mv, etc.). In 250 sec, 2068 spikes produced by cell A, 1858 spikes by cell B. c. Shared excitation. Simulated neurons similar to those above, but with no synaptic connections between them. Each independent Poisson source now produces 10-mv EPSP's, and is reduced in rate to 25/sec. An additional Poisson source with a mean rate of 150/sec delivers 10-mv EPSP's to both cells, each arriving at cell A 50 msec earlier than at cell B; i.e., each cell receives an average of 175 EPSP's per second, of which 150 (86%) are shared with the other cell. Each cell produced about 2500 spikes in 150 sec. d. Shared inhibition. Similar to previous example. Each independent excitatory input produces 10-mv EPSP's at a mean rate of 175/sec, shared channel produces 30-mv IPSP's at a mean

rate of 100/sec. Both cells depolarized by d-c input to produce spikes. About 550 spikes produced by each cell in 150 sec. e. Shared inhibition and excitation to previous cases; cell A is depolarized and receives input from shared 100/sec channel as in case 8d; cell B is excited as in example 8c. Cell A produced 732 spikes and cell B, 2600 spikes in 200 sec.




Fig. 8(a)



Fig. 8 (b)




Fig. 8(c)

1000

Number of occurrences

Ⓐ

(p) 8
 1017

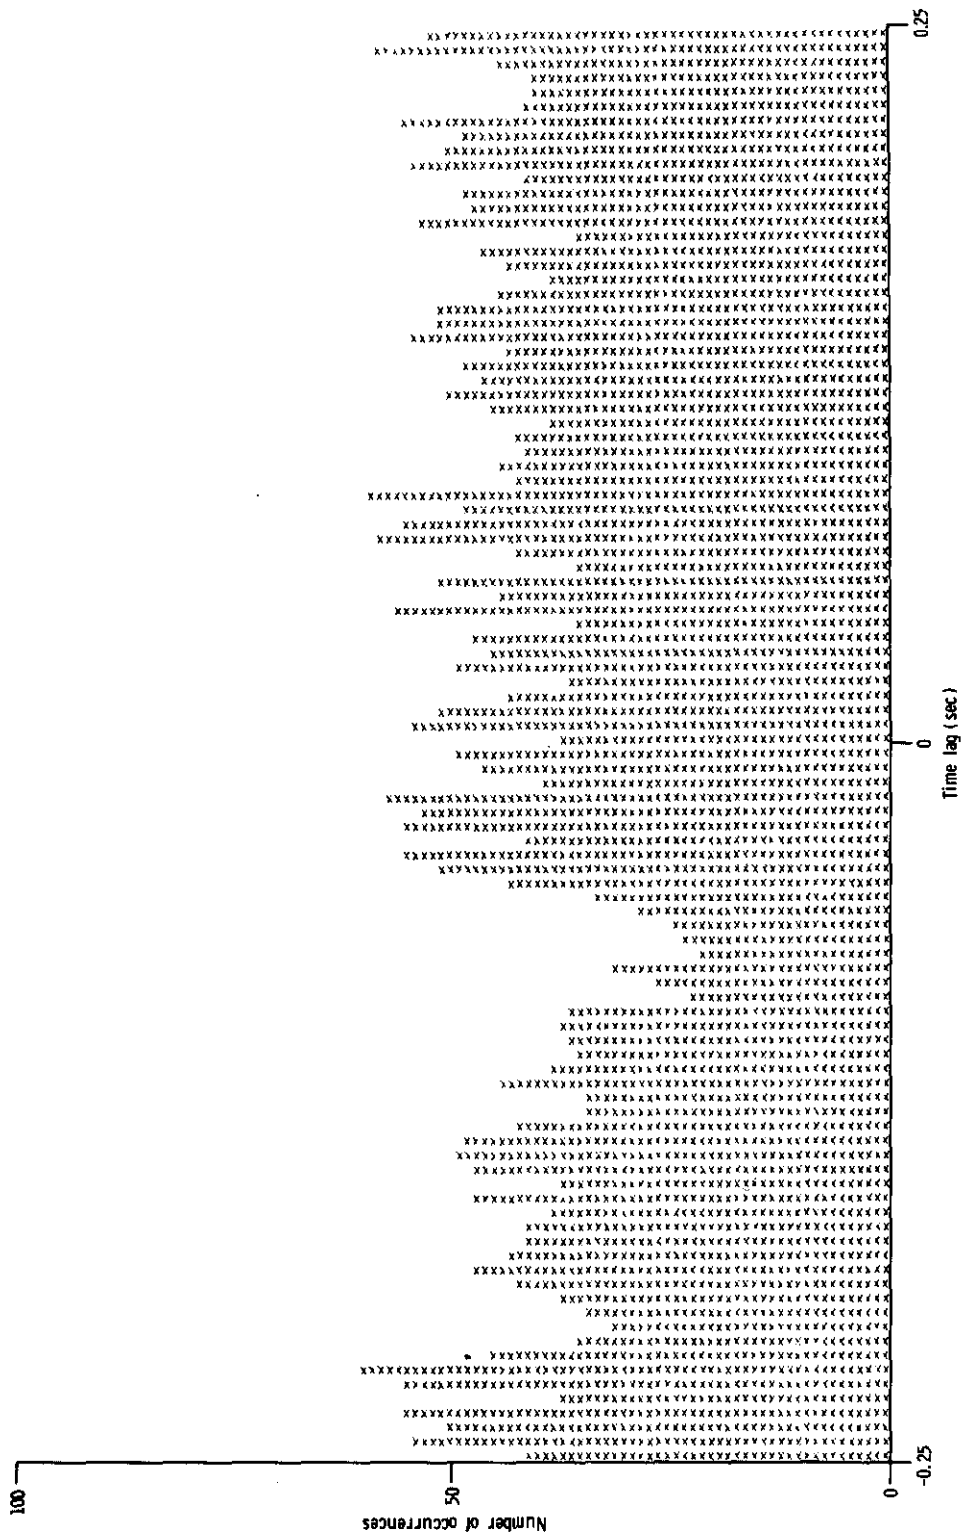


Fig. 8(e)

source might be a cell that makes excitatory synapses with cells A and B at latencies z_A and z_B , respectively. We then expect a peak in the cross-correlation function $c_{AB}(\tau)$ at or near $\tau = z_B - z_A$. The magnitude of the peak will depend on the firing rate of the source cell as well as on the neurophysiological characteristics of the observed cells.

We have examined a number of cases by means of computer simulation to obtain guidelines for the interpretation of dependence observed in physiological experiments. A few representative examples are illustrated in Fig. 8.

When networks are constructed of neurons with basically similar properties and parameters, then the following generalizations may be made on the basis of simulation studies:

- (a) Common sources of input are more difficult to detect than direct or indirect connections.
- (b) Indirect connections are more difficult to detect than direct connections.
- (c) Several different arrangements of functional interaction may lead to the same cross correlation; unique inferential conclusions concerning the anatomy and physiology of the interneuronal dependence may therefore be impossible.

In spite of these difficulties, cross correlations and their extensions can reveal a great deal about neuronal interactions. If, in addition, the organism is stimulated, a new experimental variable is provided, which can markedly

enhance the utility of spike correlation measurements, as discussed in the following section.

TWO SPIKE TRAINS IN THE PRESENCE OF STIMULATION

Functional Effects of Stimulation

Interactions between neurons may be altered by presenting a stimulus to the animal; these changes are generally reflected in the cross-correlation histogram. The effects on the cross correlation may arise (1) through changed firing rates of one or both cells, (2) through direct or indirect synaptic input to both cells from a common source (or a set of parallel sources) that respond directly to the stimulus, (3) through the effects of the stimulus on interaction pathways between the two observed cells, or (4) through any combination of these. It is shown below that the changes in the cross-correlation function can be predicted when the stimulus is repeated periodically, and when mechanisms (1) or (2), or both, are operative. If the observed changes from unstimulated to stimulated situation agree with those predicted, it may be concluded that mechanism (3) is not operative, i.e., that any interaction pathways between the two neurons are not significantly affected by the stimulus.

We now consider the case in which an identical stimulus is presented at regular intervals of duration P .^{*} These intervals are long with respect to conduction times and the duration of postsynaptic effects; i.e., the observable

^{*}Irregularly spaced stimulus presentations lead to irrelevant complications in the calculations.

effects of a given stimulus have essentially died away by the time the next stimulus is presented.

The post-stimulus-time (PST) histogram discussed above is an estimate of the firing probability of a neuron as a function of time since the onset of the stimulus. (It is clear that this is a special case of the cross-correlation function, $\zeta_{SA}(\tau)$, where "cell" S always fires at the onset of each stimulus.) It has two additive components: the relatively constant background component, which is due to firings of the observed cell that are not necessarily related to the stimulus; and departures from this level, which are due to the effects of the stimulus, whether directly or indirectly mediated. The background level may be affected by the stimulus, since ρ'_A , the mean firing rate of the cell with (repetitive) stimulus "on," may be different from ρ_A , the mean rate with stimulus "off," accordingly as the stimulus has a net excitatory or inhibitory effect on cell A.

It is clear that the observed cross-correlation function between two cells A and B will, in general, be different under "stimulus-on" and "stimulus-off" conditions. The major classes of effects possible are illustrated schematically in Fig. 9.


In the simplest situation, shown in Case a, cells A and B are independent, but both receive input from the stimulus. (Each arrow represents one or more synaptic connections,

Fig. 9.

Functional relationships between stimulus and observed neurons; schematic. Stimulus S affects neurons A and B either directly or (dotted lines) through intermediate networks X and Y. In Case a, there is no interaction between cells A and B. In Case b, there are interaction pathways I for possibly reciprocal interaction between neurons A and B. In Case c, interaction pathways I are themselves affected by the stimulus S. See text.


Case a

No interaction



Case b

Interaction



Case c

Stimulus-modulated interaction




Fig. 9

with the possible interposition of interneurons, the sketches in Fig. 9 are functional rather than anatomical diagrams.) Cells X and Y represent independent sources of synaptic activity supplying cells A and B, respectively. Cells X and Y may be influenced by the stimulus, as indicated by the broken lines. With the stimulus off, in this situation, the only contribution to the cross correlation $C_{AB}(t)$ consists of a different background component (due to rate changes) and a superimposed, non-uniform component due to shared input.

In a more complicated situation (Case b), the two cells may interact. The interaction may be from A to B, from B to A, or in both directions; it may be direct or indirect, single or multiple. There may also be shared input from sources other than the stimulus. All these possibilities are subsumed in the square box I in the figure. With the stimulus off, the background effects, interactions, and possible shared input contribute to the cross correlation (which need not be flat). When the stimulus is on, this function is affected by (1) a changed background level, (2) "first-order" shared input, represented by the arrows going directly from the stimulus S to the cells A and B, and (3) "second-order" shared input, represented by the pathways from S through A and I to B, and conversely.

A still more complicated situation is shown in Case c, in which the interaction itself is modified by

the stimulus. Such modifications, or "modulation," may result from activation of excitatory or inhibitory synaptic connections from S on interneurons in I, causing increased or reduced effectiveness in the transmission of interactions between cells A and B.

Prediction of the Cross Correlation

In the two simpler situations described above, it is possible to predict the cross correlation in the stimulus-on condition, on the basis of the following measurements: (1) the mean firing rates of both cells under both stimulated and unstimulated conditions, (2) the observed cross-correlation function with the stimulus off, and (3) the observed PST histograms for both cells.

The basic assumption involved in this prediction is that these various modifications of the cross correlation, produced by the stimulus, are sufficiently independent to be additive. This assumption breaks down, for example, when firing rates are high, so that refractory characteristics of the neurons become important.

With stimulus off, the cross-correlation histogram is composed of "background" and "interaction" components. The background component is predicted on the basis of expected random coincidences, as follows:

With a bin width of δ in the cross-correlation histogram, we define a lagged coincidence (with lag time τ) as the occurrence of an event in A at t_0 together with the

occurrence of an event in B in the interval $t_0 + \tau \pm \delta/2$.

The lag times τ correspond to midpoints of bins: $\tau_k = \pm k\delta/2$;

$k = 1, 2, \dots$. Then the expected number $N(\tau, \delta)$ of lagged coincidences in the corresponding bin, arising from a segment of A-record of length T, is given by

$$(20) \quad E[N(\tau, \delta)] = \rho_A T \int_{\tau-\delta/2}^{\tau+\delta/2} \zeta_{AB}(x) dx,$$

where $\rho_A = 1/\mu_A$ is the mean firing rate in record A.

Thus, the estimate of the theoretical density $\zeta_{AB}(\tau)$ is given by

$$(21) \quad \hat{\zeta}_{AB}(\tau) = \frac{N(\tau, \delta)}{n_A \delta},$$

where n_A is the observed number of spikes in A during T; its expected value is given by $E[n_A] = \rho_A T$, so that we have

$$(22) \quad \zeta_{AB}(\tau) = N(\tau, \delta) / (\rho_A T \delta).$$

If, however, the stationary point processes of the two cells are independent, so that an event in record A occurs without any knowledge of the sequence of events in record B, we have already pointed out that the cross correlation function is a constant: $\zeta_{AB}(\tau) = \rho_B$. Therefore, for records of finite length, the observed histogram for such independent cells has an expected value of $\rho_A \rho_B T$ in

each bin, with "noisy" fluctuations that are reduced in relative magnitude as the sample size increases. This "random," or background, level, produced through randomly occurring lagged coincidences, is the sole component of the cross-correlation function for independent cells and is a contributing component to this function for dependent cell pairs. (See Cox and Lewis (1966), pp. 246-248).

Rewriting Eq. 22, we have

$$(23) \quad N(\tau, \delta) = \rho_A T \delta \zeta_{AB}(\tau) = \rho_A \rho_B T \delta + \rho_A T \delta \xi_{AB}(\tau),$$

where $\xi_{AB}(\tau)$ represents the departure from background level, or the "interaction" component:

$$(24) \quad \xi_{AB}(\tau) \equiv \zeta_{AB}(\tau) - \rho_B.$$

With stimulus on, and with new firing rates ρ'_A and ρ'_B , the corresponding contributions to the new lagged coincidences $N'(\tau, \delta)$ are given by two terms:

$$(25) \quad G = \rho'_A \rho'_B T' \delta + \rho'_A T' \delta \xi_{AB}(\tau),$$

taking into account that the new observation time T' may differ from the previous time T .

A third component of the stimulus-on correlation histogram, which is due to shared input from the source, is obtained from the post-stimulus-time histograms. By

subtracting out the background components as in Eq. 24, we have the net PST density functions

$$(26a) \quad \xi_{SA}(\tau) = \zeta_{SA}(\tau) - \rho_A,$$

$$(26b) \quad \xi_{SB}(\tau) = \zeta_{SB}(\tau) - \rho_B,$$

which represent the departures in spike densities of cell A or B from the new mean levels at times τ after the onset of a stimulus.

Now the enhancement (or reduction) of the probability of a lagged coincidence that occurs between cells A and B can come about by an enhancement (or reduction) in delayed coincidences between S and B at a time $t + \tau$, making a contribution $\xi_{SA}(t)\xi_{SB}(t + \tau)$. We must integrate this over the entire period of the stimulus to obtain the enhancement (or reduction) per stimulus presentation of the lagged coincidence density. Thus, we evaluate the correlation integral

$$(27) \quad [\xi_{SA} \times \xi_{SB}](\tau) = \int_0^P \xi_{SA}(t)\xi_{SB}(t + \tau) dt.$$

Because the stimulus is periodic, with period P, we take the arguments modulo P in the correlation integral.*

*When the densities ξ are obtained from PST histograms, they represent averages over a bin width. If the cross-correlation histogram has the same bin structure as the

An alternative way of predicting the contribution of shared input from the stimulus to the cross correlation is to isolate those effects that are time-locked to the stimulus. This may be done as follows: one of the records (e.g., record B) is divided into equal segments, each equal to the interval P between stimulus presentations. These segments are thoroughly shuffled, so that their new order is effectively random. Then the cross correlation is recomputed between record A and shuffled record B. The shuffling has destroyed all significant time relationships between the two trains except those related to stimulus presentations. The cross-correlation histogram after shuffling, and after subtraction of the background contribution, estimates the same quantity as the correlation integral of the PST's:

$$(28) \quad \zeta'_{AB,shuf}(\tau) - \rho'_B \approx [\xi_{SA} \times \xi_{SB}](\tau).$$

In practice, shuffling need not be performed, but rather the two records are offset in time by an amount sufficient to destroy direct temporal relationships; the amount of offset is, of course, an integral multiple of the interstimulus interval P.

PST histograms, there is an imprecision in the approximation of the integral by a sum, as well as a shift in the argument τ by half a bin; the latter may be compensated by an averaging procedure.

The expected number of lagged coincidences produced by the shared input from the stimulus is given by

$$(29) \quad H = (T'_A/P)\delta [\xi_{SA} \times \xi_{SB}](\tau).$$

The total predicted number of lagged coincidences with stimulus on is given by the sum of terms G and H (Eqs. 25 and 29):

$$(30) \quad N'_{\text{pred}}(\tau, \delta) = T'\delta \left\{ \rho'_A \rho'_B + \rho'_A \xi_{AB}(\tau) + (1/P) [\xi_{SA} \times \xi_{SB}](\tau) \right\}.$$

Hence the predicted cross-correlation function is given by

$$(31) \quad \zeta'_{AB, \text{pred}}(\tau) = \xi_{AB}(\tau) + \rho'_B + (\rho'_A P)^{-1} [\xi_{SA} \times \xi_{SB}](\tau).$$

Application to Experimental Data

Comparison of the predicted and observed cross correlations enables us to distinguish among the three situations shown in Fig. 9.

If the cross-correlation function with stimulus off is flat, within statistical limits, we may conclude that there is no interaction (Fig. 9a). In this case the expected value of $\xi_{AB}(\tau)$ is zero. If the observed cross-correlation function with stimulus on does not then agree with that predicted using Eq. 31, the appropriate conclusion is that we actually have the situation of Fig. 9c, where I represents an interneuronal network that is ineffective unless potentiated by the stimulus.

If the cross correlation with stimulus off is not flat, we may conclude that there is interaction, i.e., that I is functioning. If the prediction for the stimulus-on cross correlation agrees with the actual measurement, we may conclude that we have unmodulated interaction (Fig. 9b). If there is statistically significant disagreement, we may conclude that the interaction is modulated by the stimulus (Fig. 9c). It may then be possible to distinguish among several different modes of such modulation.

The application of these techniques to the detection of interactions has been investigated through digital-computer simulations, and the results have been presented elsewhere (Perkel, 1964).

An example of this technique, drawn from experiments in the cochlear nucleus of the cat, is shown in Fig. 10. The agreement between predicted and observed cross-correlation histograms with the stimulus on indicates that the complex shape of this histogram can be completely explained by stimulus-imposed changes in firing patterns of both units.*

In summary, we may distinguish among the three possibilities as follows: (1) When there is no interaction, the cross-correlation histogram is flat with stimulus off,

*The possibility of a small excitatory influence of cell A on cell B, with a latency of about four milliseconds, however, cannot be excluded but this would have a negligible effect on the remainder of the cross-correlation histogram, which extends to half a second.

and periodic with the period of the stimulus with stimulus on; and the cross correlation can be predicted from mean rates and PST histograms. (2) When there is interaction, the cross-correlation histogram with stimulus off is not flat; with the stimulus on, it is not in general periodic, but the correlation function can still be predicted. (3) When the interaction is modulated by the stimulus, the predicted cross correlation does not agree with observation.


Fig. 10.

Effects of repeated stimulation on the cross correlation.

- a. Cross-correlation histogram between spike trains recorded from two neurons in the cochlear nucleus of the cat; unstimulated.
- b. Post-stimulus-time histogram, cell A. Stimuli consisting of brief tones were presented at 1-sec intervals.
- c. Post-stimulus-time histogram, cell B.
- d. Cross-correlation histogram under repetitive stimulation. Plotted histogram is observed; blackened points correspond to predicted cross correlation (method described in text) based on assumption of absence of potentiating effect on interactions by the stimulus. See text.

200

285



0.5

100 (sec)

0.5

Fig. 10(a)





Fig. 10(b)






Fig. 10(d)

THE PROBLEM OF NONSTATIONARITY

Basic Concepts

All the statistical measures discussed above, for single spike trains and pairs, with and without stimulation, carry the implicit assumption that the data are stationary. Specifically, this means that, in the absence of stimulation, the spike trains represent realizations of stationary point processes. In the presence of stimulation, it is assumed that the point processes are time dependent, but that the time variation is the same after each stimulus presentation; i.e., that each stimulus presentation represents a new, independent trial. Thus the stimulated cases are stationary in a larger sense.

It should be kept in mind that the phrases "stationary data" and "nonstationary data" are, strictly speaking, misnomers. The experimental data are samples, i.e., realizations over finite durations of stochastic point processes; only the (hypothetical) underlying point processes possess the properties of being stationary or nonstationary. In testing data for "stationarity," we are in fact testing whether the assumption of a stationary underlying process is a reasonable one for the body of data in question.

One of the difficulties commonly encountered in neurophysiological investigation is the fact that the behavior of a neuron under study may change significantly

during the course of observation, and therefore cannot validly be assumed to arise from a stationary process. Such changes may be exhibited in a gross way, or may be subtle and difficult to detect. The problems of detecting nonstationarity and assessing its effects are vexatious. The very meaning of stationarity depends on the context of the experiment. If, for example, a neuron undergoes a diurnal cycle of activity, a sample of a few minutes' activity may well be accepted as "stationary," whereas a sample of a few hours' activity may show marked trends, and hence be classified as nonstationary. In operational terms, therefore, it may be impossible to distinguish nonstationarity from inadequate sampling.

The most direct, straightforward, and recommended way of dealing with suspected nonstationarity is to segment the data, analyze each segment separately, and apply standard techniques for testing that the several samples were drawn from the same population. In practice, this is often impossible because of an insufficiently long sample.* A related technique, which is useful when data processing is accomplished "on line," is to observe the temporal order in which a histogram is built up. For a stationary process,

* Some classes of spike trains, which arise from certain random-walk models of neurons, correspond to renewal processes which do not have finite moments (Gerstein & Mandelbrot, 1964). A renewal process with an infinite mean would correspond to a nerve cell which has a finite probability of remaining silent indefinitely long after a spike. A

the fractional mean rate of accumulation should be uniform for all portions of the histogram (over a time period that is long with respect to any known periodic variation); any conspicuous systematic departures from this uniformity strongly suggest nonstationarity in the data.

Although the strict definition of stationarity implies that all parameters of the stochastic process are invariant with respect to a random displacement in time, detection and measurement of nonstationarity in spike trains has centered almost exclusively upon one parameter, the firing rate. In the following discussion we restrict ourselves to nonstationarities in firing rates, after pointing out the utility of statistical techniques for investigating nonstationarities in, for example, the variance of intervals, such as Bartlett's test for homogeneity of variance, (see Kendall & Stuart, 1961).

There are three principal aspects to the analysis of rate nonstationarities in spike data: The first is the detection of nonstationarity. "We may wish to test the reality of any apparent trends and this is done by testing the hypothesis of no trend" (Cox & Lewis, 1966; p. 37). To this end the reader is referred to a chapter by Cox and

nonpacemaker cell with wholly inhibitory synaptic input would remain permanently silent; some mixture of inhibitory and excitatory input could result in the cell's firing sporadically, with a long-tailed distribution of interspike intervals, which might not have a finite mean. Finite samples of such a process cannot adequately establish this possibility.

Lewis (op. cit., Chap. 3) devoted to the analysis of trends. The second aspect deals with the characterization and measurement of the variations in rate when they have been found to occur. The third aspect refers to the assessment of the effects of rate variations on other statistical measures; we consider below the effects of rate variations on the autocorrelation, the serial correlogram, and the cross correlation when the variations in rate are common to two neurons.

Measurement of Rate Variations

In spike trains exhibiting rate variation together with a relatively high degree of variability in interspike intervals, special techniques are necessary to observe the rate variations themselves, unobscured by the "local" fluctuations in interval length. One common technique for examining rate variations in a spike train is to plot each interval length as a function of time (usually taken to be the time of the second spike determining the interval) or of serial number of the interval. If the cell fires fairly regularly, such a plot will clearly reveal the structure of rate variations. If, on the other hand, the interval variance is large, trends are hidden in the "noise."

Moving-average techniques have often been used for smoothing purposes. They are typically based on either a fixed number of intervals or a fixed length of time over which an average rate is computed. A more meaningful type

of moving average would be one in which the contribution of each spike to the rate estimate is a decreasing exponential function of the time interval between the occurrence of the spike and the time referred to by the estimate. A "ratemeter" approach of this sort corresponds more closely to the response of an integrating neuron than do the more rigid moving-average techniques. The time constant for the exponential function can be chosen to correspond to the integration period of a neuron known or postulated to receive the output from the observed neuron (Segundo et al., 1966).

If trends are monotonic, estimation methods based on regression analysis are useful; these are described by Cox & Lewis (1966), Chap. 3.

Effects of Trends

One important effect of rate changes on a spike train is to increase the variation of the interspike intervals. This effect will be conspicuous, and therefore detectable in the statistics, only if this additional variability of intervals is significantly large as compared with the "intrinsic" variability of the intervals. We illustrate this with two classes of examples: a "noisy pacemaker," with an intrinsic coefficient of variation of 10%, and a Poisson process (intrinsic coefficient of variation 100%), in each of which the mean interval is a function of time.

The pacemaker results are shown in Fig. 11, in which the autocorrelation and serial correlograms are shown for a linearly accelerating train, a linearly decelerating train, and a train with sinusoidally varying mean intervals. In each of these examples, the maximum deviation of mean interval, due to rate changes, was $\pm 10\%$ of the mean interval. The effects on the autocorrelation and on the interval histograms (not shown) are not conspicuous; there is some broadening of the peaks, but this effect is apparent only upon detailed comparison with the exactly corresponding null case of no trend. Given the autocorrelations as observed experimentally, there is no reason for suspecting a trend in the data. The serial correlograms, on the other hand, clearly indicate the rate variations. An elevated serial correlogram, extending more or less uniformly out to high orders, is a

Fig. 11.

Effects of slow rate changes in a pacemaker neuron. On the left, autocorrelation histograms; on the right, serial correlograms of interspike intervals. Intervals are independently normally distributed with time-varying mean, overall mean interval in sample of 100 msec, and constant standard deviation of 10 msec. All samples are 200 sec long, approximately 2000 spikes each. (a)-(b): Accelerating pacemaker. Mean interval decreases linearly from 110 msec to 90 msec during observation period. (c)-(d): Decelerating pacemaker. Mean interval increases linearly from 90 msec to 110 msec during observation period. (e)-(f): Oscillating pacemaker. Mean interval varies sinusoidally with time, from maximum of 110 msec to minimum of 90 msec, with a 2-sec period.




Fig. 11(a)




Fig. 11(b)



Fig. 11(c)




Fig. 11 (d)




Fig. 11(e)




Fig. 11 (r)

specific indicator of monotonic trend. It is to be noted that the effects of monotonic acceleration are indistinguishable from those of monotonic deceleration. The undamped oscillatory nature of the third serial correlogram shown is due to the imposed constancy of period of rate variation; if that period had varied during the observation, the oscillations in the serial correlogram would have exhibited damping. An ingenious method is described by Firth (1966) for separating the effects of trend on the serial correlogram from its "inherent" features. His technique, involving successive differences, is in essence a form of analysis of variance, and is applied to cells that fire at extremely regular intervals.

The time-varying Poisson process, on the other hand, requires a rate variation of $\pm 50\%$ before some effects are noticeable, and only at $\pm 70\%$ rate variation do the effects become conspicuous (Fig. 12). The only effect on the autocorrelation of a monotonic trend is to increase its level; the shape remains flat, as in the absence of trend. The predicted asymptotic level for the autocorrelation depends only on the mean observed interval (Eq. 12); therefore, the observed autocorrelation histogram, together with this predicted level, can indicate a trend. For the monotonically time-varying Poisson processes illustrated (Fig. 12), this is the only conspicuous effect of the large rate variations, since the corresponding serial correlograms depart only slightly from zero. At these

Fig. 12.

Effects of slow rate changes in a nonpacemaker neuron.

On the left, autocorrelation histograms; on the right, serial correlograms of interspike intervals. Intervals are generated by a Poisson process with time-varying rate parameter; period of observation is 200 sec, with approximately 2000 spikes each sample. Reference level in autocorrelation is predicted asymptotic value (see text, Eq. 12).

(a)-(b): Deceleration, + 50%. Mean interval increases linearly with time from 50 msec to 150 msec. (c)-(d):

Deceleration, + 70%. Mean interval increases linearly with time from 30 msec to 170 msec. (e)-(f): Oscillation,

+ 50%. Mean interval varies sinusoidally from minimum of 50 msec to maximum of 150 msec, with period of 0.75 msec.

(g)-(h): Oscillation, + 70%. Mean interval varies sinusoidally from minimum of 30 msec to maximum of 170 msec, 0.75-sec period.




Fig. 12(a)




Fig. 12(b)




Fig. 12(c)





Fig. 12(d)






Fig. 12(f)




Fig. 12(g)




Fig. 12(b)

intensities of cyclic rate variation, oscillations are visible in the autocorrelation, but not in the serial correlogram; this is in contrast with the opposite situation observed in the case of a pacemaker with weak oscillations in rate. A monotonically declining autocorrelation histogram, such as the early portions of Figs. 12c and 12d, is strongly suggestive of rather severe rate changes in the data, and some detailed statistical features of the spike train may be masked or distorted thereby.

Rate changes shared by two otherwise independent neurons may, if sufficiently pronounced, be revealed in two-cell comparisons. It can be shown, for two Poisson processes with time-varying parameters, that the maximum fractional expected departure of the cross correlation from its "null" level is less than the square of the maximum fractional rate variation. If, for example, the shared rate changes of the two processes vary within 20% of their respective means, then the cross correlation will depart by at most 4% from its predicted value for independent stationary processes based upon observed firing rates (Eq. 15). The actual amount of departure depends upon the precise nature of the temporal variation of the rate parameters of the Poisson processes; details will be presented elsewhere.

These results are illustrated in Fig. 13, in which cross-correlation histograms are shown for pairs of Poisson processes subjected to monotonic and oscillatory rate

Fig. 13.

Effects of slow rate changes on cross correlations.
Cross-correlation histograms between pairs of trains drawn from Poisson processes with time-varying rate parameters. Each member of a pair of trains has identical time variation of mean rate, but each train is generated independently. Reference level in cross correlation is predicted expected level for stationary independent processes, based upon observed mean firing rates (see text, Eq. 15). Rate variations same as in Fig. 12. a. Deceleration, $\pm 50\%$. b. Deceleration, $\pm 70\%$. c. Oscillation, $\pm 50\%$. d. Oscillation, $\pm 70\%$.




Fig. 13(a)




Fig. 13(b)




Fig. 13(c)




Fig. 13(d)

changes. The effect is always an elevation in level of the cross-correlation histogram in the neighborhood of its origin. Linear trends result in a uniformly elevated cross correlation, which remains flat. Oscillating trends give rise to an oscillating cross-correlation histogram, if the range of the latter is great enough.

Because the effects of nonstationarity on the cross correlation are of second order, the statistical indications of interaction between two cells (discussed in an earlier section), are not severely affected by rate changes in one or both of the cells, even if they are fairly severe. Inasmuch as stationary conditions are often difficult to maintain in an experiment, this fact is an encouragement to the attempt to elucidate functional interconnections of neurons through statistical comparison of their spike trains.

REFERENCES

- Amassian, V. E., J. Macy, Jr., H. J. Waller, H. S. Leader, and M. Swift (1962), "Transformation of Afferent Activity at the Cuneate Nucleus," pp. 235-254 in Information Processing in the Nervous System, Physiol. Sci., 22nd Congress, Leiden, 1962. Vol. 3, Excerpta Medica Found.
- Bartlett, M. S. (1963), "The Spectral Analysis of Point Processes," J. Royal Statistical Soc. B, Vol. 25, pp. 264-296.
- Bullock, T. H. (1957), "Neuronal Integrative Mechanisms," in Recent Advances in Invertebrate Physiology, B. T. Scheer, (ed.), University of Oregon Press, Eugene, Oregon.
- Burns, B. D., and G. K. Smith (1962), "Transmission of Information in the Unanesthetized Cat's Isolated Forebrain," J. Physiol., Vol. 164, pp. 238-251.
- Cooley, J. W., and J. W. Tukey (1965), "An Algorithm for the Machine Calculation of Complex Fourier Series," J. Math. Comp., Vol. 19, pp. 297-301.
- Cox, D. R. (1962), Renewal Theory, John Wiley & Sons, Inc., New York.
- Cox, D. R., and P. A. W. Lewis (1966), The Statistical Analysis of a Series of Events, John Wiley & Sons, Inc., New York.
- Cox, D. R., and H. D. Miller (1965), The Theory of Stochastic Processes, John Wiley & Sons, Inc., New York.
- Firth, D. R. (1966), "Interspike Interval Fluctuations in the Crayfish Stretch Receptor," Biophys. J., Vol. 6, pp. 201-215.
- Geisler, C. D. and J. M. Goldberg (1966), "A Stochastic Model of the Repetitive Activity of Neurons," Biophys. J., Vol. 6, pp. 53-69.
- Gerstein, G. L. (1960), "Analysis of Firing Patterns in Single Neurons," Science, Vol. 131, pp. 1811-1812.
- Gerstein, G. L., and W. A. Clark (1964), "Simultaneous Studies of Firing Patterns in Several Neurons," Science, Vol. 143, pp. 1325-1327.
- Gerstein, G. L., and N. Y.-S. Kiang (1960), "An Approach to the Quantitative Analysis of Electrophysiological Data from Single Neurons," Biophys. J., Vol. 1, pp. 15-28.

- Gerstein, G. L., and B. Mandelbrot (1964), "Random Walk Models for the Spike Activity of a Single Neuron," Biophys. J., Vol. 4, pp. 41-68.
- Goldberg, J. M., H. O. Adrian, and F. D. Smith (1964), "Response of Neurons of the Superior Olivary Complex of the Cat to Acoustic Stimuli of Long Duration," J. Neurophysiol., Vol. 27, pp. 706-749.
- Hagiwara, S. (1949), "On the Fluctuation of the Interval of the Rhythmic Excitation. I. The Efferent Impulse of the Human Motor Unit During the Voluntary Contraction," Rept. Physiograph, Sci. Inst. Tokyo Univ., Vol. 3, pp. 19-24.
- Hagiwara, S. (1954), "Analysis of Interval Fluctuation of the Sensory Nerve Impulse," Japan. J. Physiol., Vol. 4, pp. 234-240.
- Huggins, W. H. (1957), "Signal-flow Graphs and Random Signals," Proc. IRE, Vol. 45 (Part 1), pp. 74-86.
- Junge, D., and G. P. Moore (1966), "Interspike-interval Fluctuations in Aplysia Pacemaker Neurons," Biophys. J., Vol. 6, pp. 411-434.
- Kendall, M. G., and A. Stuart (1961), The Advanced Theory of Statistics, Vol. 2, 2nd ed., Charles Griffin, London.
- Lamarre, Y., and J. P. Raynauld (1965), "Rhythmic Firing in the Spontaneous Activity of Centrally Located Neurons. A Method of Analysis," Electroencephalog. Clin. Neurophysiol., Vol. 18, 87-90.
- Lewis, P. A. W. (1966), "A Computer Program for the Statistical Analysis of Series of Events," IBM Systems J., Vol. 5, No. 4 (in the Press).
- McFadden, J. A. (1962), "On the Lengths of Intervals in a Stationary Point Process," J. Roy. Statist. Soc. Ser. B, Vol. 24, pp. 364-382.
- McGill, W. J. (1963), "Stochastic Latency Mechanisms," pp. 309-360 in Handbook of Mathematical Psychology, Vol. 1, R. D. Luce, R. R. Bush, and E. Galanter (eds.), John Wiley & Sons, Inc., New York.
- Moore, G. P., D. H. Perkel, and J. P. Segundo (1966), "Statistical Analysis and Functional Interpretation of Neuronal Spike Data," Ann. Rev. Physiol., Vol. 28, pp. 493-522.

Perkel, D. H. (1964), Detection of Functional
Among Neurons: A Technique Using Repetitive Present
of Stimuli, The RAND Corporation, RM-4234-NIH.

Perkel, D. H. (1965), "Applications of a Digital Computer
Simulation of a Neural Network",
and Cybernetic Systems, M. Maxfield, A. Callahan,
L. J. Fogel (eds.), Spartan Books, Washington, D.C.

Poggio, G. F., and L. J. Viernstein (1964), "Time Series
Analysis of Impulse Sequences of Thalamic Somatic Sensory
Neurons," J. Neurophysiol., Vol. 27, pp. 517-545.

Rodieck, R. W., N. Y.-S. Kiang, and G. L. Gerstein (1962),
"Some Quantitative Methods for the Study of Spontaneous
Activity of Single Neurons," Biophys. J., Vol. 2,
pp. 351-368.

Segundo, J. P., D. H. Perkel, and G. P. Moore (1966),
"Spike Probability in Neurones: Influence of Temporal
Structure in the Train of Synaptic Events," Kybernetik,
Vol. 3, pp. 67-82.

Simon, W. (1965), "The Real-time Sorting of Neuro-electric
Action Potentials in Multiple Unit Studies," Electro-
encephalog. Clin. Neurophysiol., Vol. 18, pp. 192-195.

Smith, D. R., and G. K. Smith (1965), "A Statistical
Analysis of the Continual Activity of Single Cortical
Neurones in the Cat Unanaesthetized Isolated Forebrain,"
Biophys. J., Vol. 5, pp. 47-74.

Watson, G. S., and M. R. Leadbetter (1964a), "Hazard
Analysis I," Biometrika, Vol. 51, pp. 175-184.

Watson, G. S., and M. R. Leadbetter (1964b), "Hazard
Analysis II," Sankhya (Indian J. Statistics), Vol. 26,
pp. 101-116.

Weiss, T. F. (1964), "A Model for Firing Patterns of Auditory
Nerve Fibers," Tech. Report No. 418, Research Laboratory
of Electronics, Massachusetts Institute of Technology,
Cambridge, Massachusetts.

Wilson, D. M. (1964), "Relative Refractoriness and Patterned
Discharge of Locust Flight Motor Neurons," J. Exp. Biol.,
Vol. 40, pp. 191-205.

Wyman, R. (1965), "Probabilistic Characterization of Simul-
taneous Nerve Impulse Sequences Controlling Dipteran
Flight," Biophys. J., Vol. 5, pp. 447-471.